

(19) World Intellectual Property Organization
International Bureau(43) International Publication Date
13 September 2001 (13.09.2001)

PCT

(10) International Publication Number
WO 01/67803 A1(51) International Patent Classification: **H04Q 11/04.**
H04L 12/56

[GB/GB]: The Old Police House, Park Road, Grundisburgh, Woodbridge, Suffolk IP13 6TP (GB).

(21) International Application Number: **PCT/GB01/00972**(74) Agent: **WILSON, Peter, David**; BT Group Legal Services, Intellectual Property Dept., Holborn Centre 8th floor, 120 Holborn, London EC1N 2TE (GB).

(22) International Filing Date: 6 March 2001 (06.03.2001)

(25) Filing Language: English

(81) Designated States (national): CA, US.

(26) Publication Language: English

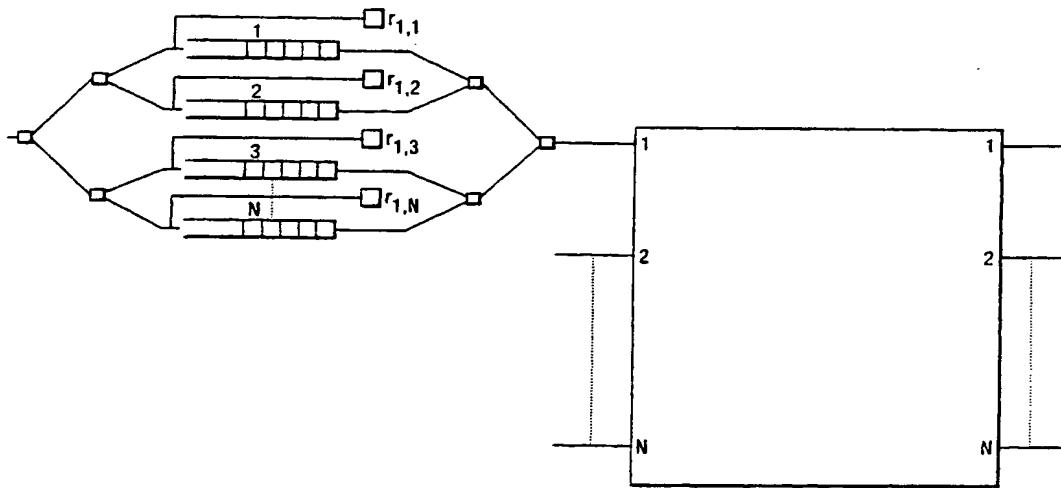
(84) Designated States (regional): European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR).

(30) Priority Data:
0006084.8 10 March 2000 (10.03.2000) GB**Published:**

- with international search report
- before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

(72) Inventor; and

(75) Inventor/Applicant (for US only): **HILL, Alan, Michael**(54) Title: **PACKET SWITCHING**

(57) **Abstract:** A method of allocating switch requests within a packet switch, the method comprising the steps of establishing switch request data at each input port; processing the switch request data for each input port to generate request data for each input port-output port pairing; comparing the number of requests from each input port and to each output port with the maximum request capacity of each input port and each output port; allocating all requests for those input-output pairs where the total number of requests is less than or equal to the maximum request capacity of each input port and each output port; reducing the number of requests for those input-output pairs where the total number of requests is greater than the maximum request capacity of each input port and each output port such that the number of requests is less than or equal to the maximum request capacity of each input port and each output port; and allocating the remaining requests. Packets may be switched from an input port to a specified output port in accordance with the allocations obtained with the above method.

WO 01/67803 A1

PACKET SWITCHING

This invention relates to packet switching (or cell switching), in particular methods for allocating requests for switching from one of the inputs of a packet switch to one of the outputs of the packet switch.

Input-buffered cell switches and packet routers are potentially the highest possible bandwidth switches for any given fabric and memory technologies, but such devices require scheduling algorithms to resolve input and output contentions. Two approaches to packet or cell scheduling exist (see, for example, 5 A Hung *et al*, "ATM input-buffered switches with the guaranteed-rate property," A Hung *et al*, Proc. IEEE ISCC '98, Athens, Jul 1998, pp 331-335). The first approach applies at the connection-level, where bandwidth guarantees are required. A suitable algorithm must satisfy two conditions for this; firstly it must 10 ensure no overbooking for all of the input ports and the output ports, and 15 secondly the fabric arbitration problem must be solved by allocating all the requests for time slots in the frame.

Fabric arbitration has to date been proposed by means of the Slepian-Duguid approach and Paull's theorem for rearrangeably non-blocking, circuit-switched Clos networks (see Chapter 3, J Y Hui, *Switching and traffic theory for integrated broadband networks*, Kluwer Academic Press, 1990). This connection-level algorithm can be summarised as firstly ensuring no overbooking and 20 secondly performing fabric arbitration by means of circuit-switching path-search algorithms. It has been assumed that this algorithmic approach could only be applied at the connection level, because of its large computational complexity. For 25 this reason, proposals for scheduling of connectionless, best-efforts packets or cells employ various matching algorithms, many related to the "marriage" problem (see D Gale and L S Shapley, "College admissions and the stability of marriage," 30 *Mathematical Monthly*, 69, 9-15 (1962) and D Gusfield and RW Irving, *The Stable Marriage Problem: Structure and Algorithms*, MIT Press, 1989) in which the input-output connections for each time slot or phase of the switch are handled independently, i.e. a frame of time slots (and hence phases) is not employed. Although such algorithms for choosing a set of conflict-free connections between 35 inputs and outputs for each time slot, which are based on maximum size and

maximum weight bipartite graph matching algorithms, can achieve 100% throughput (N McKeown *et al.*, "Achieving 100% throughput in an input-queued switch," Proc. IEEE Infocom '96, March 1996, vol.3, pp.296-302) they are also impractically slow, requiring running times of complexity $O(N^3 \log N)$ for every time 5 slot (R E Tarjan, "Data structures and network algorithms," Society for Industrial and Applied Mathematics, Pennsylvania, Nov. 1983).

Iterative, heuristic, parallel algorithms such as iSLIP are known, which reduce the computing complexity (i.e. time required to compute a solution) for best-efforts packets or cells (N McKeown *et al.*, "The Tiny Tera: a packet switch 10 core," IEEE Micro Jan/Feb 1997, pp 26-33). The iSLIP algorithm is guaranteed to converge in at most N iterations, and simulations suggest on average in fewer than $\log_2 N$ iterations. Since no guarantees are needed, this and similar algorithms currently represent the preferred scheduling technique for connectionless data at the cell level in input-buffered cell switches and packet routers with large numbers 15 of ports (e.g. $N \geq 10$). The iSLIP algorithm is applied to the Tiny Tera packet switch core, which employs Virtual Output Queueing (VOQ), in which each input port has a separate FIFO (First In, First Out) queue for each output, i.e. N^2 FIFOs for an $N \times N$ switch. If we assume that each FIFO queue stores at least a number 20 of cells equal to the average cell latency L , and that each cell is a 53-byte ATM cell, then the total input FIFO queue hardware count is $O(424LN^2)$. With each element capable of clocking out 424f bits per frame, this is a complexity product of $O((424)^f fLN^2)$, which is a very large complexity. Fortunately, by employing a single queue in the form of RAM in each port, acting as N virtual queues, the hardware count can be reduced to $O(424LN)$, and with parallel readout reducing 25 the number of steps per frame to just f , the overall complexity product can be reduced to $O(424fLN)$. Table 1 gives the hardware and "computing" steps for these queues to provide f cells within a frame.

For unicast packets the iSLIP algorithm converges in at most N iterations, where N is the number of input and output ports. On average the algorithm 30 converges in fewer than $\log_2 N$ iterations. The physical hardware implementation employs N round-robin grant arbiters for the output ports and N identical accept arbiters for the input ports. Each arbiter has N input ports and N output ports, making N^2 links altogether. The total amount of hardware depends on the precise

construction of the round-robin arbiters. N McKeown *et al.*, *op cit*, employ a priority encoder to identify the nearest request from the port closest to a pre-determined highest-priority port (see Figure 1). The priority encoder reduces the number of links down to $\log_2 N$ parallel links, in order to change the pointer if required. The $\log_2 N$ parallel links are then expanded back up to N links again through a decoder. Details of the hardware complexity of the arbiters are given in N McKeown, Scheduling Algorithms for Input-Queued Cell Switches, PhD Thesis, University of California, Berkeley, 1995. The growth rate for the complete scheduler is $O(N^4)$, each arbiter being $O(N^3)$. For a 32x32 cell switch (which is the size of the Tiny Tera switch), 421,408 2-input gates are required. This may be quite acceptable for such a small switch, but the $O(N^4)$ growth rate is extremely large.

In order to minimise the overall hardware and computing complexity, the best structure for constructing the encoder is a binary tree, which requires $O(2N)$ elements (for large N) and only $\log_2 N$ steps per iteration, whilst the decoder needs only $O(N)$ elements. Pipelining cannot be employed to reduce to one step per iteration, because the pointers cannot be up-dated until the single-bit requests have passed through both sets of arbiters to the decision register. The total hardware and computing complexities are given below in Table 1. The hardware complexity now grows as $O(N^4)$ rather than $O(N^4)$, due to the binary tree encoder and decoder.

	Hardware Count	Computing Steps per Frame	Hardware Computing Complexity Product
Input RAM queues	$424LN$	f	$424fLN$
Average Convergence	$O(6N^2)$	$O(4f\log_2 N(1 + \log_2 N))$	$O(24fN^2\log_2 N(1 + \log_2 N))$
Guaranteed Convergence	$O(6N^3)$	$O(4fN(1 + \log_2 N))$	$O(24fN^3(1 + \log_2 N))$

Table 1. Hardware and computing complexities of the iSLIP algorithm for scheduling f packets per port in a frame of f time slots.

The overall hardware computing complexity product $O(24fN^3(1 - \log_2 N))$ of the iSLIP algorithm for scheduling f packets per port would be no less than that of the maximum size and weight matching algorithms of N McKeown, *et al* "Achieving 100% throughput in an input-queued switch," Proc. IEEE Infocom '96, March 5 1996, vol.3, pp.296-302., if convergence must be guaranteed. There is a reduction to $O(24fN^2\log_2 N(1 - \log_2 N))$ for the average number of computing steps. The major benefit of the iSLIP algorithm is its parallel nature, which allows the number of computing steps to be traded against hardware complexity, thus reducing computing times by a factor N^2 at the expense of increasing the 10 hardware by the same factor. It is interesting to note that hardware quantities for the input RAM queues far exceed those needed for the scheduling electronics.

According to a first aspect of the invention there is provided a method of a method of allocating switch requests within a packet switch, the method comprising the steps of

- 15 (a) establishing switch request data at each input port;
- (b) processing the switch request data for each input port to generate request data for each input port-output port pairing;
- (c) comparing the number of requests from each input port and to each output port with the maximum request capacity of each input port and each 20 output port;
- (d) allocating all requests for those input-output pairs where the total number of requests is less than or equal to the maximum request capacity of each input port and each output port;
- (e) reducing the number of requests for those input-output pairs where the 25 total number of requests is greater than the maximum request capacity of each input port and each output port such that the number of requests is less than or equal to the maximum request capacity of each input port and each output port; and
- (f) allocating the remaining requests.

30 According to a second aspect of the invention there is provided a method of allocating switch requests within a packet switch, the method comprising the steps of;

- (a) establishing switch request data at each input port;

(b) processing the switch request data for each input port to generate request data for each input port-output port pairing;

(c) allocating a first switch request from each of the input port-output port pairing request data, the requests being allocated only if the maximum request capacity of the respective output port has not been reached; and

5 (d) allocating further switch requests by the iterative application of step (c) until the maximum request capacity of each output port has been reached.

The present invention additionally provides a method of allocating switch requests within a packet switch, the method comprising the steps of:

10 (a) establishing switch request data at each input port;

(b) processing the switch request data for each input port to generate request data for each input port-output port pairing;

(c) identifying a first switch request from each of the input port-output port pairing request data;

15 (d) identifying further switch requests by the iterative application of step (c) until all of the switch request data has been identified;

(e) subject to the maximum request capacity of each input port and each output port, allocating all of the identified switch requests; and

(f) reserving unallocated switch requests for use in the next phase of switch

20 request allocation.

The invention will now be described with reference to the following figures in which:

Figure 1 is a schematic depiction of a known arrangement for allocating switch requests;

25 Figure 2 is a schematic depiction of an apparatus for counting switch requests according to the present invention;

Figure 3 is a schematic depiction of a second apparatus for counting switch requests according to the present invention;

30 Figure 4 is a schematic depiction of an apparatus for counting the switch requests for each output port of a packet switch according to the present invention;

Figure 5 is a schematic depiction of an apparatus for counting the switch requests for each output port of a packet switch according to an alternative embodiment of the present invention;

Figure 6 is a graph comparing the performance of the iSLIP algorithm with 5 that of the present invention;

Figure 7 is a graph showing the performance ratio of the iSLIP algorithm to that of the present invention;

Figure 8 is a second graph comparing the performance of the iSLIP algorithm with that of the present invention; and

10 Figure 9 is a second graph showing the performance ratio of the iSLIP algorithm to that of the present invention.

As the scheduling of best-effort, connectionless cells within an input-buffered switch, router or network is being considered, each of the input ports 15 could be assumed to have a FIFO queue, each of which is destined for a different output port (i.e. virtual output queueing - VOQ). Although the flows are best-effort, we wish to schedule them on a frame-by-frame basis. However, there is no pre-reservation of time slots within this frame, thus a number f of cells are queued 20 at each input port, in f time slots, and are being scheduled according to their output port destinations in such a way as to avoid contention. A particular cell or packet should be able to be transmitted across the switch fabric during any one of f time slots. Before performing fabric arbitration to ensure that there is no output port contention in each time slot, we must first make sure that there is no overbooking of the input and output ports within the frame.

25 If the total number of cells Nf (where N is the number of input ports and f is the number of time slots in a frame) to be switched across a cell or packet switch are to be computed together on a frame-by-frame basis, by means of a path-searching algorithm for a 3-stage circuit switch, then every cell could be represented as a port on the circuit switch. The number of computing steps 30 needed to ensure no overbooking then depends on the amount of hardware that is acceptable. If $O(fN\log_2(fN))$ components are used, then $O(fN)$ computing steps are needed. The number of computing steps can be reduced to $O(\log_2^2(fN))$ if more hardware is acceptable, i.e. $O(fN\log_2^2(fN))$, using a Batcher sorting network, but

5 this hardware quantity may be too large to be acceptable. However, to represent every cell as a port on a circuit switch is an over-restrictive constraint in a cell switch. In fact, in a cell switch, it is only necessary to ensure that the number of cells destined for each of the N output ports does not exceed the number of cells or time-slots in the frame as there is no requirement to exit the output port in any specific time-slot.

The present invention concerns a method of for ensuring no overbooking of the input and/or output ports. An $N \times N$ request matrix R is defined, whose elements r_{ij} represent the number of cells in input port i destined for output port j .

10 The two conditions that ensure no overbooking are simply:

$$\sum_{i=1}^N r_{ij} = f \quad \text{for all } j \quad \text{and} \quad \sum_{j=1}^N r_{ij} = f \quad \text{for all } i$$

15 In practice, cells from more than f time slots in each input port could be considered in this procedure, if cells destined for overbooked ports have to be discarded. Discarded cells could either be lost completely, or continue to be queued for later attempts.

20 In order to establish the number of request matrix elements, N^2 counters are established, one for each queue. Figure 2 shows a schematic depiction of a possible input port arrangement for counting the request matrix elements, r_{ij} . Each of the N input ports 10 to a switch fabric 20 has N FIFO queues 11, N counters 12 and N switches 13 in order to direct cell requests to the appropriate counter 12. Assuming just f cells in each port are counted within the request counter 12. Assuming just f cells in each port are counted within the request matrix, each counter then requires $\log_2(f+1)$ counting stages, requiring $O(N^2 \log_2 f)$ counter elements altogether. If it is assumed that the individual cell destination requests are input to these counters as single bits, there will be a maximum of $O(f)$ computing steps required of any counter, giving an overall hardware computing complexity product for the counters of $O(fN^2 \log_2 f)$. Figure 2 shows that we also require $O(N)$ switches in each input port to direct the cell requests to the correct counter, i.e. $O(N^2)$ in total. The speed of these switches must be sufficient that, within one frame of f slots, $f \log_2 N$ bits can be routed. The overall complexity product for the switches is therefore $fN^2 \log_2 N$.

The method of queueing cell requests can be refined further. In Figure 2, each port has N FIFO queues, each capable of buffering up to f cells. This requires $O(fN^2\log_2N)$ buffer elements, each capable of toggling \log_2N times per frame, i.e. a complexity product of $O(f^2N^2(\log_2N)^2)$. Fortunately, by employing a single queue

5 in the form of RAM in each port, acting as N virtual queues, the hardware count can be reduced to $O(fN\log_2N)$, with the same number of steps per frame, requiring a complexity product of $O(f^2N(\log_2N)^2)$ overall. Since a particular cell stored in RAM will be allocated to any of the f time slots within the frame, re-ordering may now also be required in an output queue if it is desired to preserve the cell order

10 between input and output ports. This is now in essence the same as traditional time-slot interchanging in time-shared circuit switches. However, it would be possible to preserve the cell order of a virtual input queue by allocating time slots in time order to the cells destined for the same output port. Even with efficient buffering in a single queue within each input port, cell buffering is the most

15 complex function in a switch, requiring the largest quantity of the fastest electronics. The cells here are just the request cells, containing only output port destination addresses (and possibly input port addresses as well as other parameter values). Much greater buffering complexity is needed for the actual cells or packets carrying all the header information and payload.

20 Figure 3 shows an improved arrangement for counting the request matrix elements. A serial input stream of cell requests is converted to a parallel word, which is then transmitted over a parallel bus 31. Each line of the parallel bus is connected via gates 32 to sections of the RAM 33, each RAM section 34 holding an individual cell request of \log_2N bits. Each RAM section can both read and

25 write cell requests from and to the parallel bus. As the cell requests are written into RAM sections 34, they are also decoded into single-bit requests by the decoder 35 and transmitted to the array of counters 36. Each input port requires an array 36 having N counters, so the requirement for an $N \times N$ switch is N^2 . The overall complexity product, which was previously dominated by the RAM queues,

30 is now reduced by a factor \log_2N , which for large N could be an order of magnitude. This reduction is in terms of RAM access speeds, rather than quantity of buffer memory.

Once all of the matrix elements have been counted, the next step of the method of the present invention is to add the request matrix elements r_{ij} to form the sum of the requests for each output port of the switch. If more than f requests are taken into account, then the sum of the requests from each input port must also be calculated. Figure 4 shows an array of N^2 counters 41, each containing the number of requests for switching from a given input port to a given output port, e.g. counter [1,2] holds requests for switching from the first input port to the second output port and counter [1,N] holds requests for switching from the first input port to the N th output port. The outputs from the counters 41 feed into an array of N adders 42, with each adder corresponding to one of the output ports of the switch fabric. Thus, the outputs of the counters which hold requests for switching to the first output port all connect to the input of the adder 42 which is associated with the first output port of the switch fabric. The counts may be represented as $\log_2 f$ -wide words, each of which must be switched successively to the adder circuitry 43 by an associated switch array 44. For conventional adder constructions the software and hardware complexities are no greater than for counting the individual request matrix elements.

The third step of the method of the present invention is to compare the summations for each output port with f , which is the maximum number of cells that can be sent from each input port and to each output port in each frame. If any row or column of the request matrix exceeds f , the number of requests must be reduced to no more than f . One method of achieving this is to reduce the number of allowed requests to a number proportional to the actual number of requests, i.e.

25

$$r'_{ij} = \frac{f}{\sum r_{ij}} \bullet r_{ij}$$

This step is efficient when there is a heavy concentration of requests on one, or a few, input or output ports. However, if the traffic is distributed uniformly from inputs to outputs, such that each element has a small number of requests, then the method is less efficient. In an alternative embodiment of the present

invention, the requests are allocated in the following manner. Allocation of requests must be performed fairly, so that every request has a chance of being granted within the frame, while satisfying the "no overbooking" condition, to keep the average delay seen by individual requests low. Furthermore, every virtual 5 output queue must have a chance of being granted a request within the frame, to prevent starvation. To fulfil these requirements, each r_{ij} from all input ports must be granted at least one request, if they ask for one or more requests. as one port, can not be granted a large number of requests while other ports are granted none.

The tasks of summing requests and then reducing the number of requests 10 as necessary are now replaced by a single mechanism which iteratively counts up the requests by granting one at a time to all r_{ij} counters, until the sum of the requests destined for any output port equals the number f of slots in the frame. At that point no more requests can be granted to that output. On the first iteration, all request matrices, r_{ij} , with one or more requests for a given output 15 port j are granted one of these requests, so that up to N requests may be granted in the first iteration (assuming here that $f \geq N$). Each of the non-zero request matrix, r_{ij} , counters is now decremented by one, (i.e. a 1 is subtracted from each $\log_2 f$ -wide word) ready for the next iteration. Meanwhile, the successful requests granted to each output port, which are indicated as single bits in parallel from 20 each r_{ij} counter, are summed, for example by an adder with parallel input ports, converting the individual granted request bits into a word of length $\log_2(2N)$ bits.

Figure 5 shows a circuit which can be used to implement the alternative embodiment of the invention. The circuit includes an array of N^2 request matrices 41 (as shown in Figure 4 and described above) and N adding elements 51. The 25 adding elements 51 each comprise a parallel bit adder 52 which receives in parallel the single bit requests from each associated request matrix 41 and creates a $\log_2(2N)$ -wide word. The output of each parallel bit adder 52 is connected to a respective $\log_2(2f)$ -wide adder 53.

There are N such parallel adders 52, each of which requires $4N$ binary 30 adders. The number of steps required in each iteration to obtain the sum of the requests is $\log_2 N \cdot \log_2(2N)/2$. Those r_{ij} counters that have a second request for a given output port have a second single bit summed through the parallel adders in a second iteration of the process described above. Altogether there could be as

many as f iterations, if only one input port has requests destined for any one of the output ports. At the other extreme, for a uniform traffic distribution where a given output receives cells from different inputs, there could be just one iteration required. Because there is no buffering within the parallel adders, pipelining can be 5 employed to sum each of the iterations. The maximum number of steps required for a maximum of f iterations is $(f + \log_2 N \cdot \log_2(2N)/2)$. The complexities are summarised in Table 4.

At the output of each parallel adder 52 there is a temporal succession of up to $f \log_2(2N)$ -wide words. The successive words must also be added 10 sequentially by the adder 53 to obtain the overall total number of requests granted to an output port. Since this number cannot exceed f (i.e. the number of time slots in the frame) this can be done using a $\log_2(2f)$ -wide adder. The adder construction in Figure 5 employs $O(\log_2(f))$ half adders. Alternatively, a conventional adder construction employing $\log_2(2f)$ full adders could be used.

15 The succession of $f, \log_2(2N)$ -wide words can be stored in a buffer of size $f \log_2(2N)$, if necessary. These can then clocked out sequentially at a suitable rate via switches for the following $\log_2(2f)$ -wide adder. The $\log_2(2f)$ -wide adder 53 calculates the overall total number of requests granted to all input ports destined 20 for a particular output port (and this total must not exceed f). On the iteration that drives the overall total above f , the counting must be stopped for that output port. Each request matrix, r , counter whose cell requests are destined for that output port must be advised and keep a record of the number of iterations it took part in. It will only be allowed a number of requests equal to the number of iterations for which the overall count of requests for a given output port is $\leq f$. If 25 the overall total on the previous iteration is less than f , then up to N additional requests may be allocated up to the total of f (when $f \geq N$). This can simply be done by examining the status of each request matrix, r , counter in turn in a maximum of N steps. When $f < N$, then up to f additional requests may be allocated. This may also take a maximum of N steps, by examining the status of 30 each request matrix counter in turn, depending on the locations of the requests distributed around the N counters. Each of the request matrix counters now knows how many requests have been allocated to it. To maintain fairness between the input ports destined for a particular output port, a pointer may be

used, so that additional requests can be allocated preferentially to different input ports in different frames. (Of course, a specific pattern of requests could be such that the same input ports are in fact allocated additional requests in different frames). There are many ways in which a pointer could be used to allocate 5 additional requests, including existing round-robin techniques. The simplest way would be to cycle the pointer continuously around the input ports, allocating one request at a time to any requesting port, stopping when the required number of additional requests up to N has been granted. The next frame that subsequently needs to allocate additional requests then begins from that pointer position.

10 It would be possible to reduce the N steps required to allocate the additional requests to $O(\log_2 N)$ steps, by using an (N, N) concentrator to pack requests next to each other, so that only the required number up to N of those concentrated requests can be gated or switched through to the parallel bit adder. The concentrator construction should preferably be such that the relative 15 positions of the requests are preserved at the output of the concentrator, like that in T Szymanski, "Design principles for practical self-routing non-blocking switching networks with $O(N \log N)$ bit-complexity," IEEE Trans. On Computers, vol.46, no.10, 1057-1069 (1997).. The gating or switching of the concentrator outputs could be achieved by means of a modified $\log_2 N$ -stage decoder, which 20 not only produces, for example, a logic "1" on the appropriate numbered decoder output port (controlling the concentrator output port), to represent the last of up to N ports to be allowed through to the adder, but also propagates through its $\log_2 N$ stages a logic "1" to all decoder outputs (and controlled concentrator outputs) that lie above the last of up to N ports. In this way all decoder output 25 ports above and including the decoded one provide enable bits ("1"s) to the concentrator output ports that they control, and all the decoder ports below the decoded one provide disable bits ("0"s) to the concentrator output ports that they control. Because the pointer could be in any of the N counter positions, and we wish preferably to allow through requests starting from the pointer position, the 30 process of allocating additional requests could be split into two steps.

In the first, only the requests including and below the pointer position are sent to the concentrator for gating or switching through to the adder. Disabling of the requests above the pointer position can be performed by a similarly modified

decoder. In the second part, only those requests lying above the pointer position are allowed through to the concentrator. The overall hardware complexity required for the concentrator and decoder using this construction is $O(3N^2\log_2 N)$, and the number of computing steps using Szymanski's concentrator construction is 5 $O(8\log_2 N)$.

In a further alternative embodiment of the present invention, the counting and summing of requests could be performed as described above above, except that the counting and summation may continue beyond the iteration that drives the overall total of requests above f for a particular output port, so that the 10 complete total of requests can be counted. The $\log_2(2f)$ -wide adder would now need to be a $\log_2(2fN)$ -wide adder, with a hardware count of $N\log_2(2Nf).\log_2(4Nf)/2$; computing steps $O(f\log_2(2Nf))$; and overall complexity 15 $O(Nf\log_2^2(2Nf).\log_2(4Nf).2)$. Furthermore, it could also be possible for the counters to hold cell request counts from previous frames, and to add these to the counts of each new frame. This would allow running totals, or perhaps weighted averages, to be used covering longer time-slots than a single frame, if cell queue lengths employed are longer than a frame length. The number of cell acceptances between input and output ports could then be calculated on perhaps a fairer basis related to longer-term flows between input and output ports.

Indeed, whether the actual queue lengths employed are longer than the 20 frame length or not, i.e. even if unsuccessful cells are discarded within each frame, it may be advantageous for the number of cell acceptances within each frame to be related to such a longer-term measure of traffic flow requests between ports, rather than just the cell requests within the frame itself. Of course 25 the number of cell acceptances could also be related to a combination of longer-term flow and "within frame" requests.

Once the number of requests r allocated between each input and output port within a frame is known, the individual cell requests buffered in the RAM 30 port must be identified. This can be achieved by re-running the individual cell queues must be identified. This can be achieved by re-running the individual cell requests out of the RAM queues, through the decoder switches and re-counting them in the request matrix counters. This time, each counter is set at its allocated total, and can be decremented by each cell request bit that it receives. While the 35 total, and can be decremented by each cell request bit that it receives. While the counter is still above zero, a single bit (e.g. a "1") is sent back to the RAM queue counter is still above zero, a single bit (e.g. a "1") is sent back to the RAM queue

to the cell position corresponding to the current request bit, signifying cell acceptance. When a cell request bit decrements the counter to zero, and for all subsequent request bits, a single bit (e.g. a "0") is sent back to the appropriate RAM queue cell position to signify non-acceptance of that particular cell request.

5 The status (accepted or rejected) of all cell requests stored in RAM queues is now established. The hardware and computing complexities are the same as for the first count of matrix elements r_{ij} , except that an additional $N(N + f)$ switches are needed. The additional number of steps required through these is f . All accepted cell requests are now ready to have their time slots calculated for transmission
10 across the switch fabric (fabric arbitration).

Excluding the necessary cell or packet buffering in RAM queues, for which there is no essential difference, the hardware/computing complexity product of the method of the present invention is $O(\log_2 N)$ smaller than iSLIP for all hardware items. Taking the worst values across all hardware items, again excluding cell
15 buffering, this actually translates into fewer computing steps than iSLIP by the same factor $O(\log_2 N)$, but at the expense of greater hardware quantities by the same factor $O(\log_2 N)$ for some of the hardware items. Nevertheless, the hardware quantities required for the method of the present invention are much smaller than the cell buffering hardware.

20 Figure 6 shows the number of computing steps for the iSLIP algorithm with average convergence (solid line) and the method of the present invention (dashed line), for a small switch with $N = 32$ input and output ports, as a function of the number of time slots f under consideration. Figure 7 shows the ratio of computing steps for the two algorithms, which reaches a minimum of 0.32 for $f = 32$ time
25 slots. Although this means that the method of the present invention is more than three times as fast an algorithm as iSLIP, it does require cells to be buffered for 32 time slots to achieve this minimum ratio. However, the ratio is around 1.3 for all numbers of time-slots from 8 upwards, so any desired cell latency could be chosen. The benefits of a frame-based algorithm become more significant for
30 switches with more ports N . Figures 8 and 9 show the equivalent graphs as Figures 6 and 7 for the case where $N = 256$ ports. Here, the method of the present invention takes only 0.195 times the iSLIP computing time at minimum, requiring an optimum $f = 192$ time slots. Once again, the practical number of time

slots can be anything from 64 upwards, yet still provide around a 5-fold speed advantage. Thus significant computing time reductions are achievable even if the number of time slots f employed is made equal to the number of ports N .

The overall hardware computing complexity product of the method of the present invention is of order N times smaller than a maximum weight matching algorithm. In comparison, the iSLIP algorithm having average convergence is only $O(N/\log_2 N)$ smaller.

CLAIMS

1. A method of allocating switch requests within a packet switch, the method comprising the steps of

- 5 (a) establishing switch request data at each input port;
- (b) processing the switch request data for each input port to generate request data for each input port-output port pairing;
- (c) comparing the number of requests from each input port and to each output port with the maximum request capacity of each input port and each 10 output port; and
- (d) allocating all requests for those input-output pairs where the total number of requests is less than or equal to the maximum request capacity of each input port and each output port;
- (e) reducing the number of requests for those input-output pairs where 15 the total number of requests is greater than the maximum request capacity of each input port and each output port such that the number of requests is less than or equal to the maximum request capacity of each input port and each output port; and
- (f) allocating the remaining requests.

20

2. A method of allocating switch requests within a packet switch, the method comprising the steps of:

- (a) establishing switch request data at each input port;
- (b) processing the switch request data for each input port to generate 25 request data for each input port-output port pairing;
- (c) allocating a first switch request from each of the input port-output port pairing request data, the requests being allocated only if the maximum request capacity of the respective output port has not been reached; and
- (d) allocating further switch requests by the iterative application of step 30 (c) until the maximum request capacity of each output port has been reached.

3. A method of allocating switch requests within a packet switch, the method comprising the steps of:

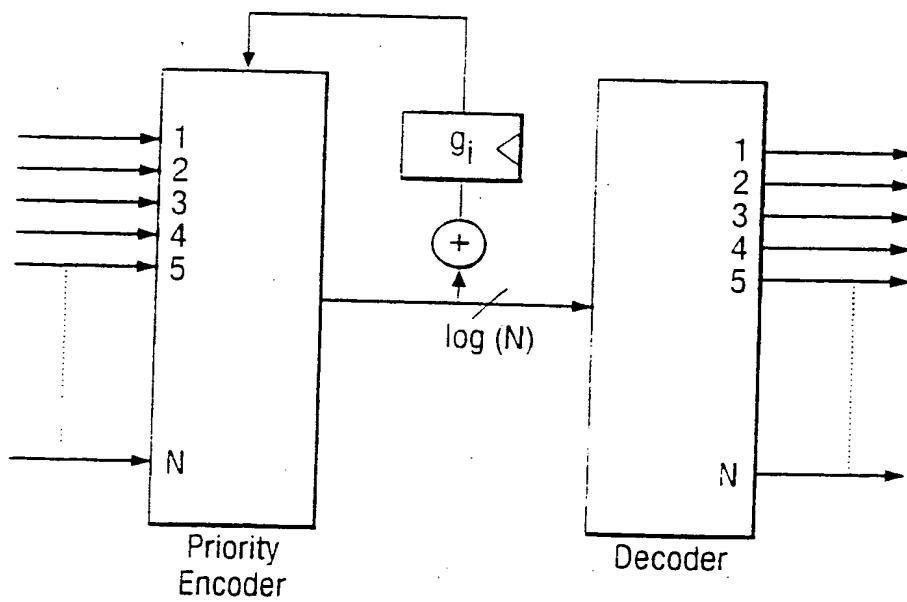
- (a) establishing switch request data at each input port;
- (b) processing the switch request data for each input port to generate request data for each input port-output port pairing;
- (c) identifying a first switch request from each of the input port-output port pairing request data;
- (d) identifying further switch requests by the iterative application of step (c) until all of the switch request data has been identified;
- (e) subject to the maximum request capacity of each input port and each output port, allocating all of the identified switch requests; and
- (f) reserving unallocated switch requests for use in the next phase of switch request allocation.

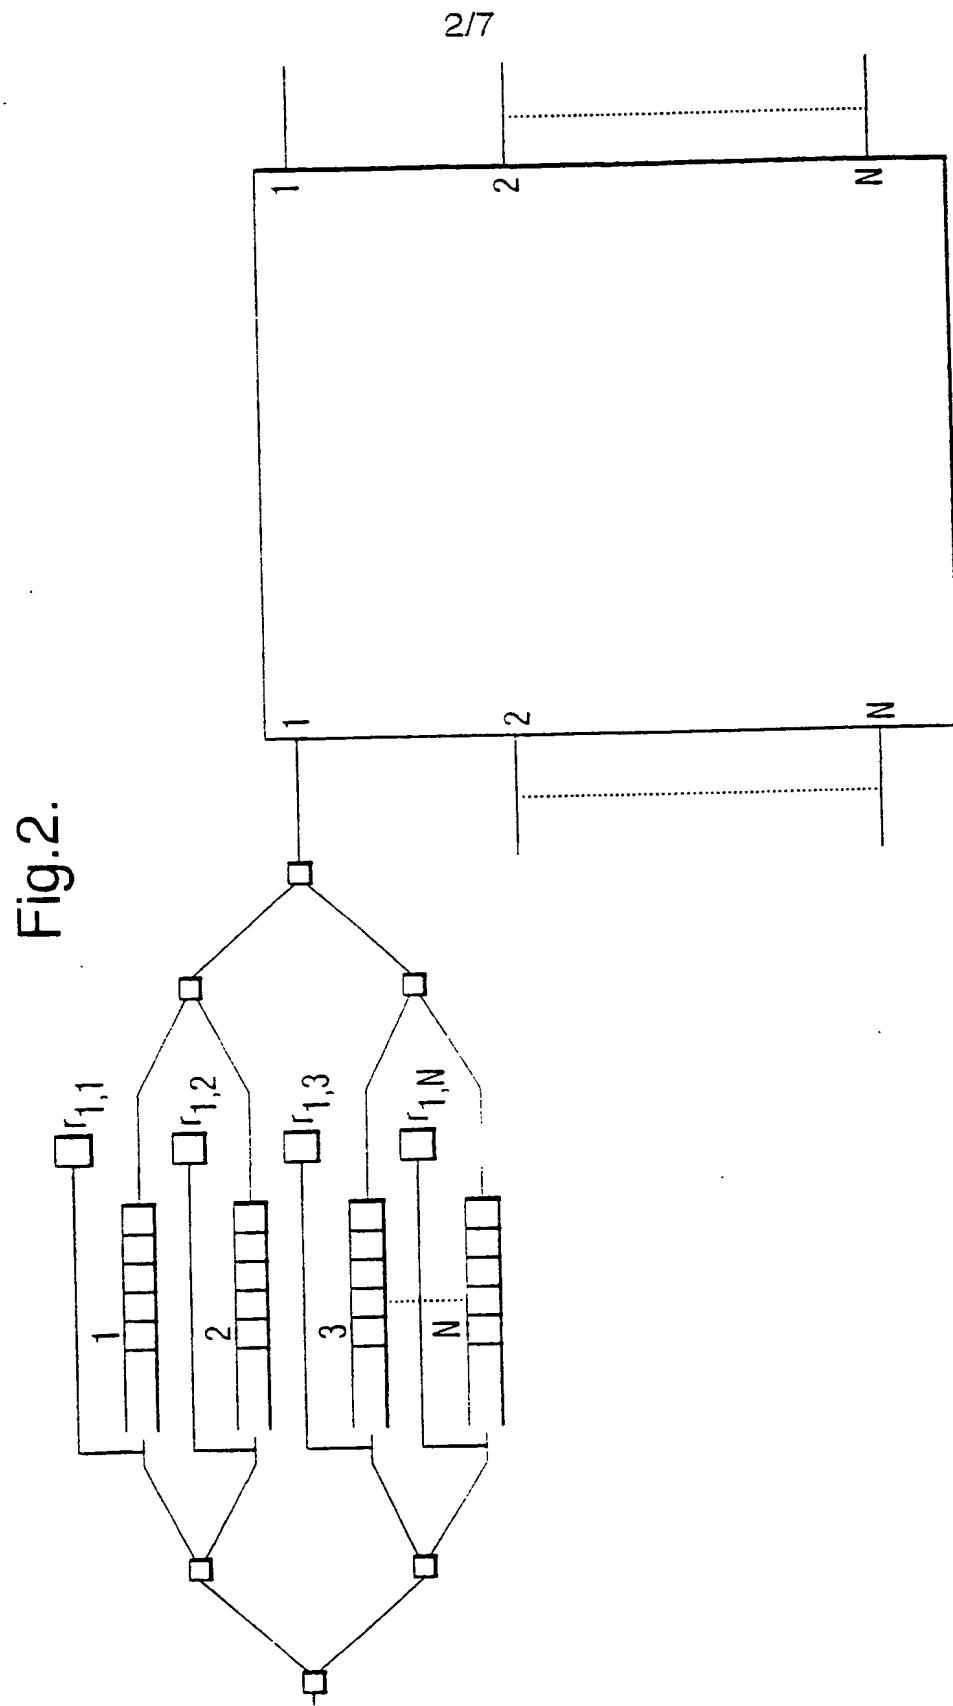
4. A method of packet switching wherein the input port-output port routing is allocated according to the method of any of claims 1-3 and the packets are switched on the basis of the allocated routing.

5. A packet switch in which switch requests the input port-output port routing is allocated in accordance with the method of any of claims 1 to 3.

20 6. A packet switch according to claim 5, wherein packets are switched from an input port to a specified output port in accordance with the allocated routing.

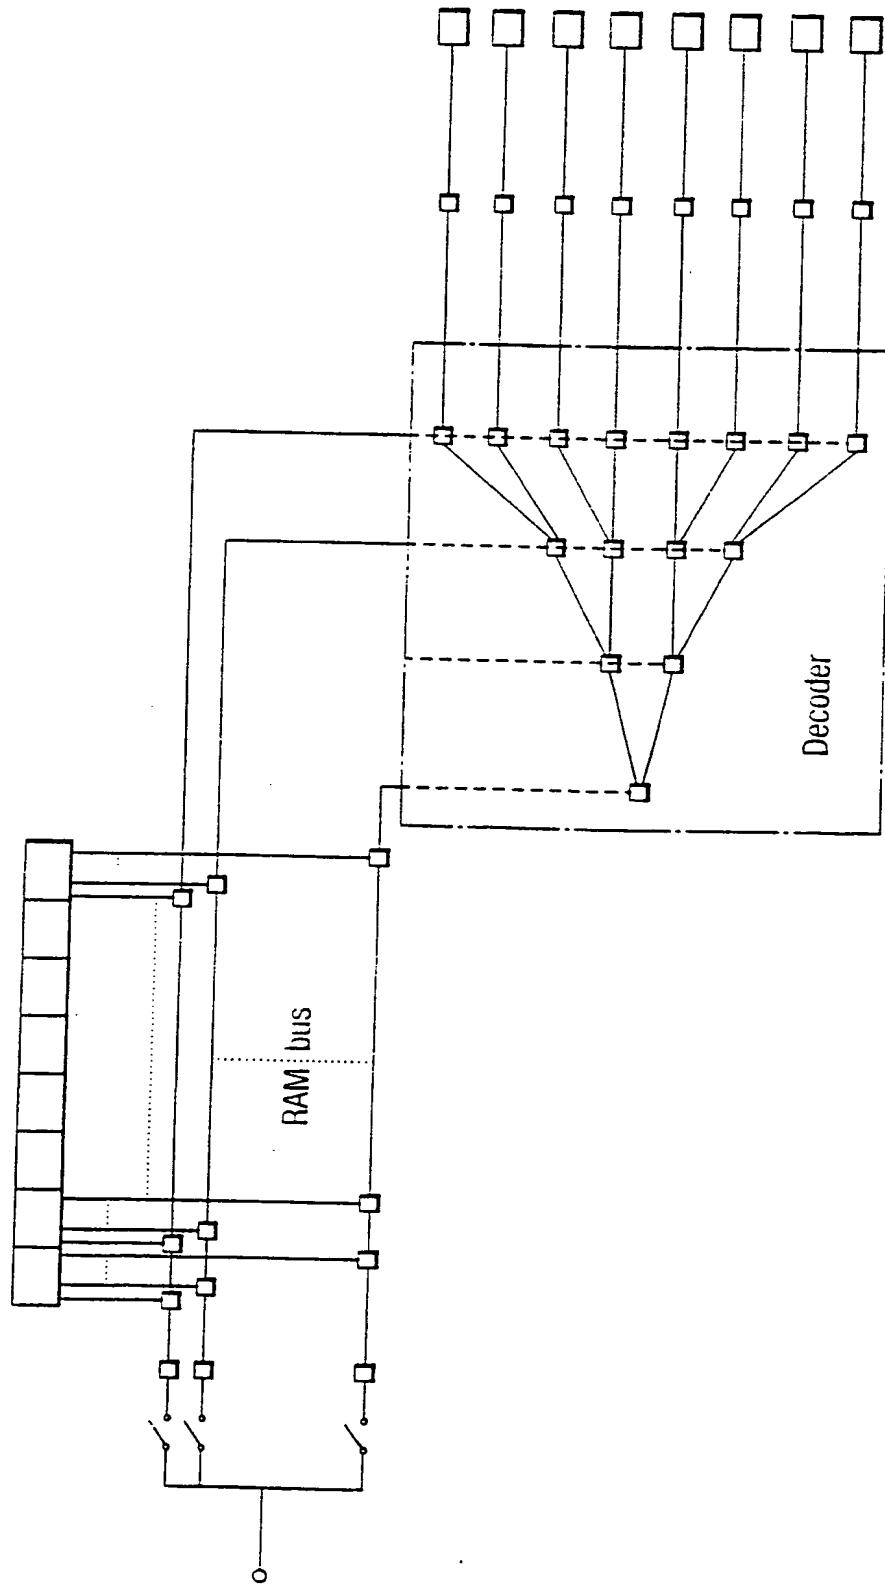
Fig.1.



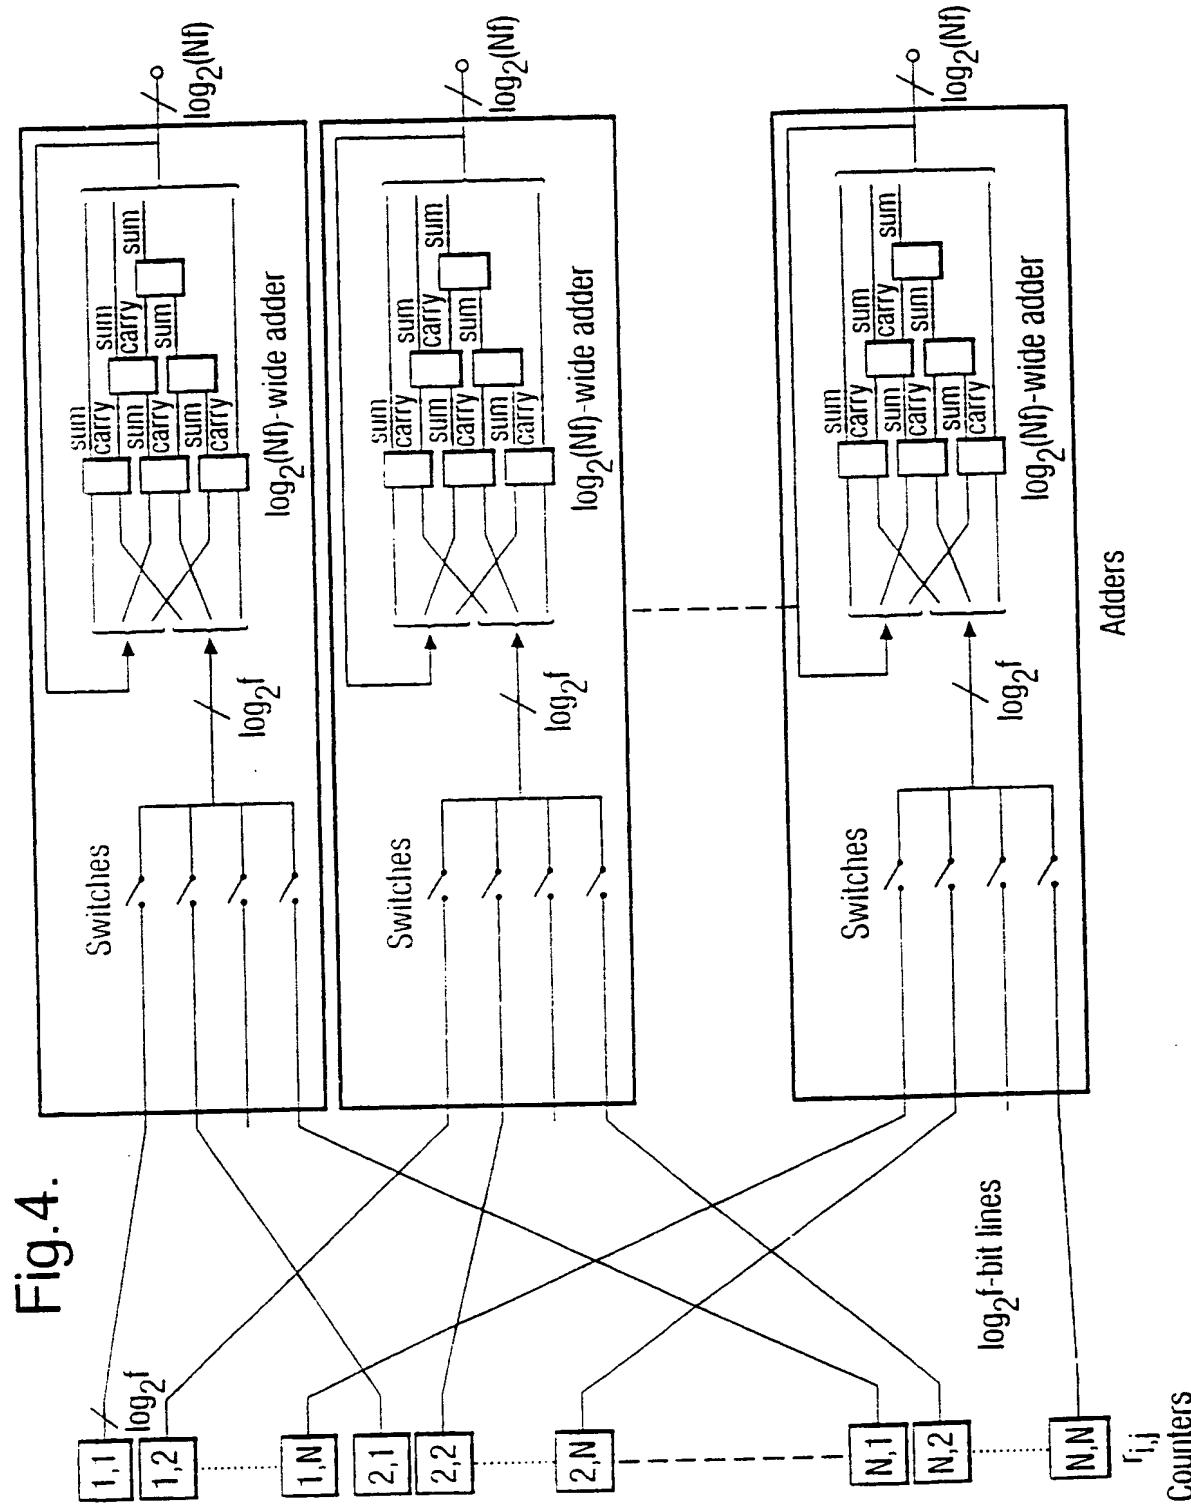


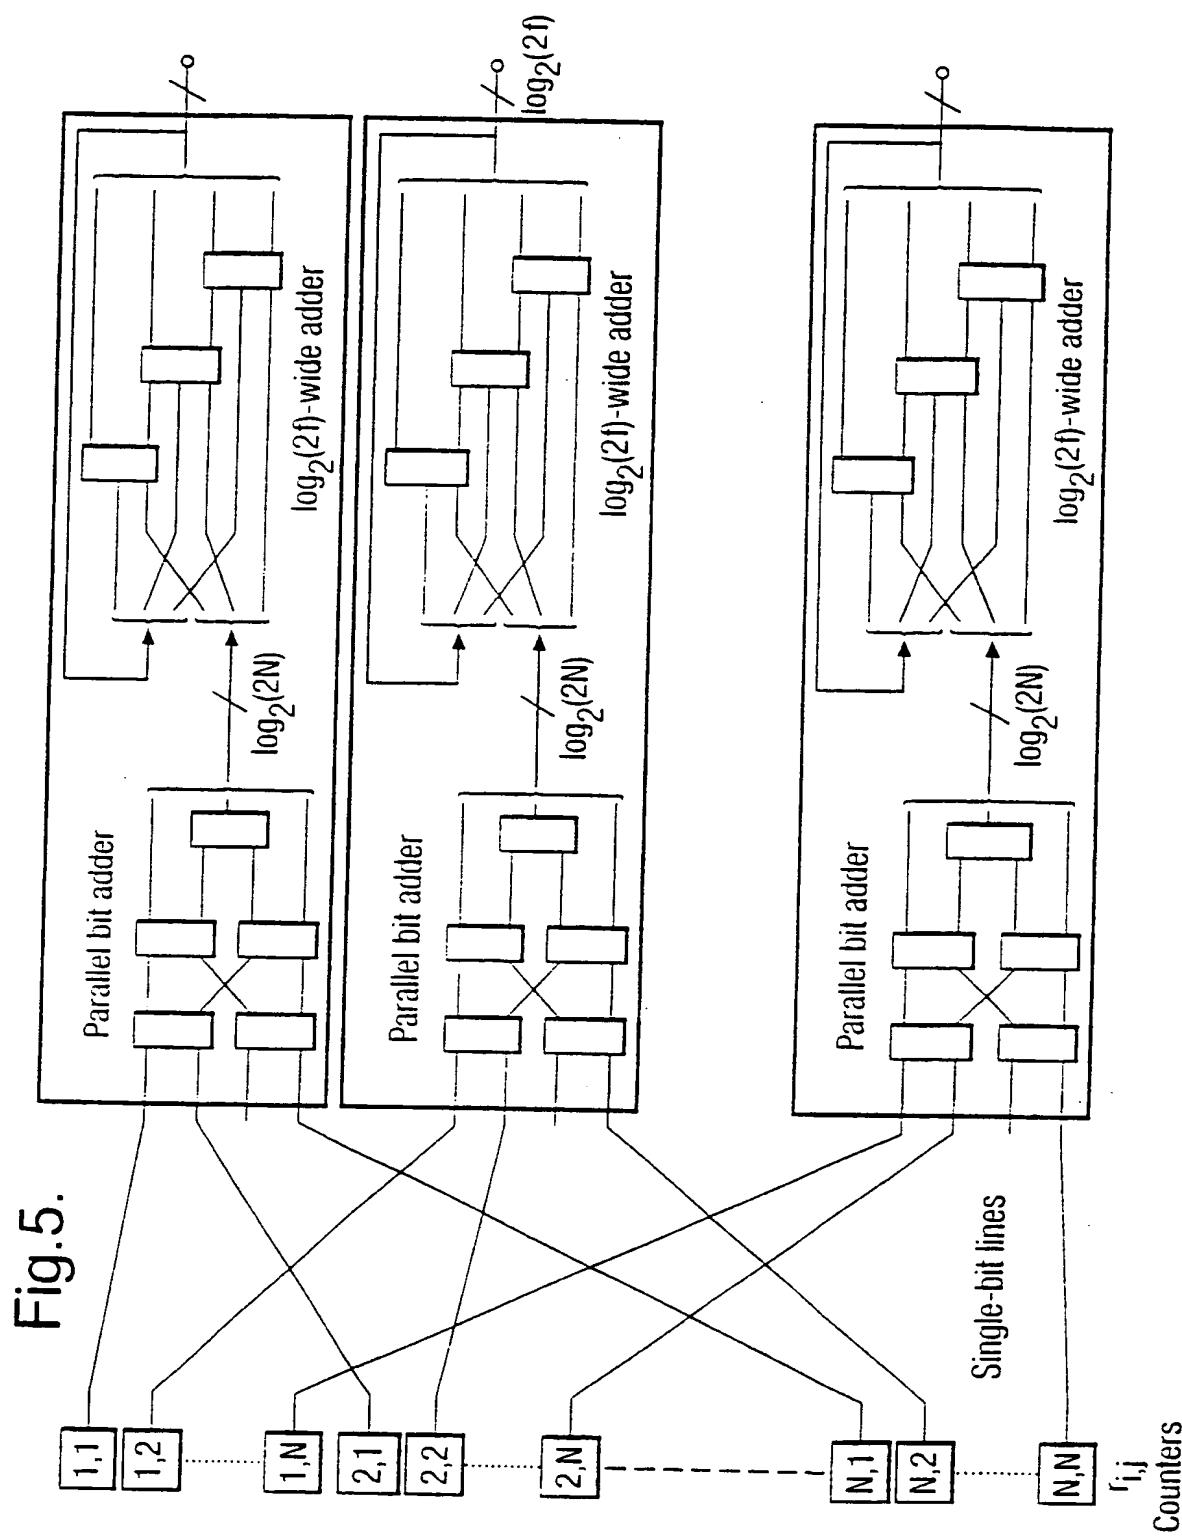
3/7

Fig.3.



4/7





6/7

Fig.6.

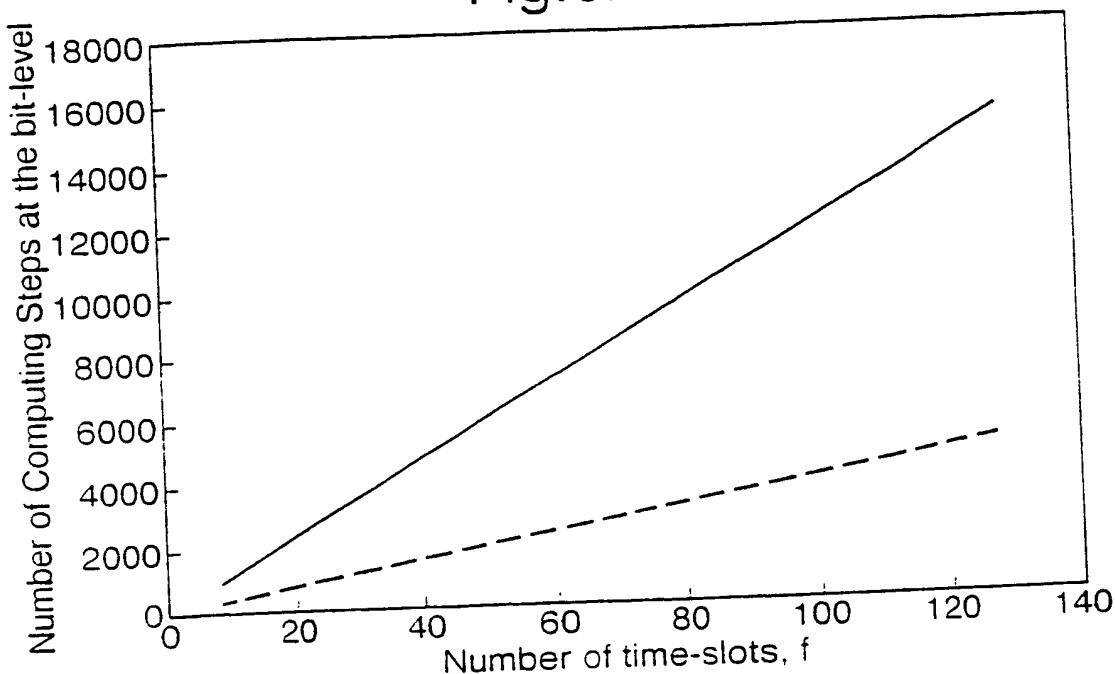
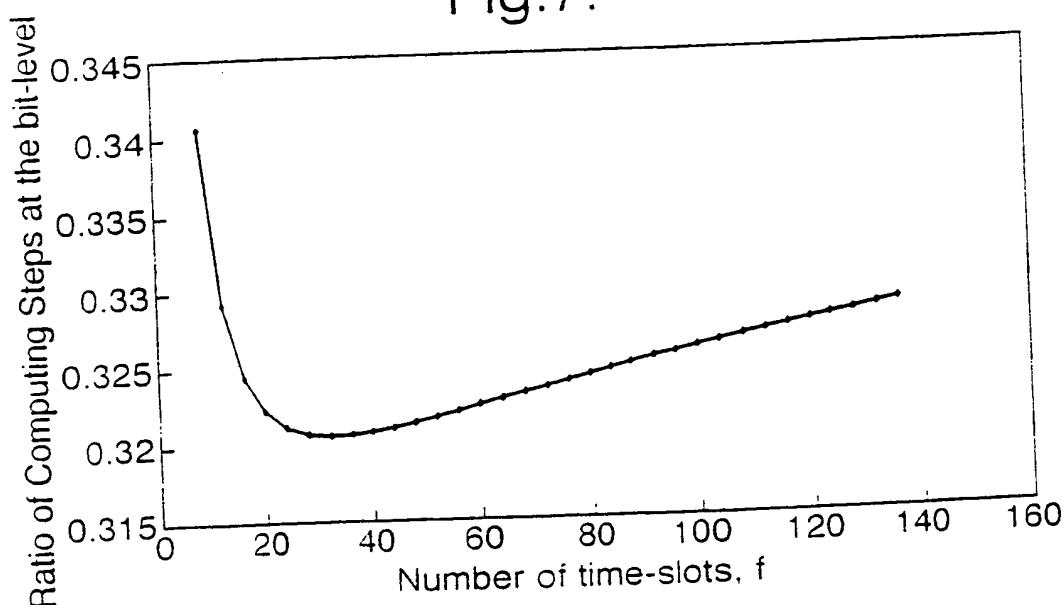


Fig.7.



7/7

Fig.8.

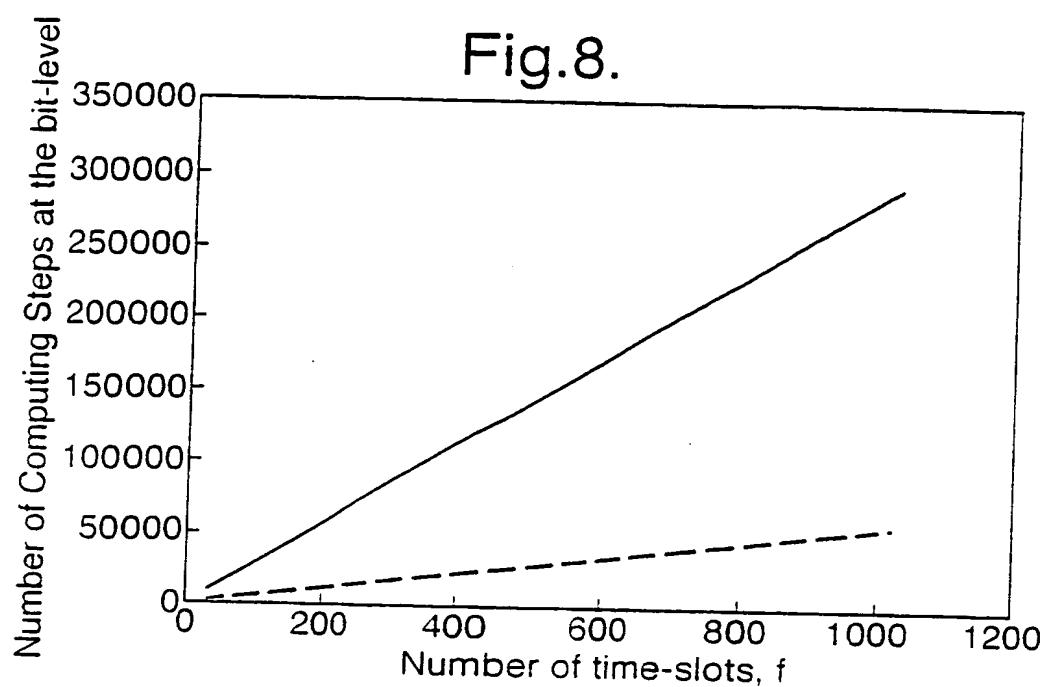
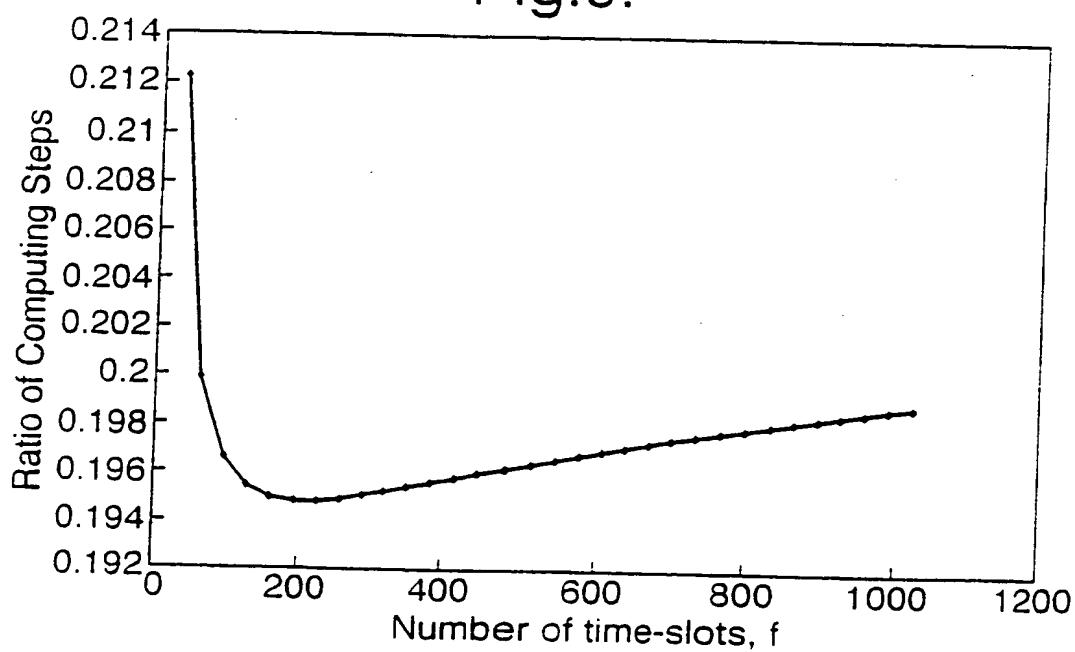


Fig.9.



INTERNATIONAL SEARCH REPORT

International Application No

PCT/GB 01/00972

A. CLASSIFICATION OF SUBJECT MATTER
 IPC 7 H04Q11/04 H04L12/56

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 7 H04Q H04L

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

EPO-Internal

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	<p>CHEN W -T ET AL: "A HIGH PERFORMANCE CELL SCHEDULING ALGORITHM IN BROADBAND MULTICASTSWITCHING SYSTEMS" GLOBAL TELECOMMUNICATIONS CONFERENCE (GLOBECOM), US, NEW YORK, IEEE, page 170-174 XP000737529 ISBN: 0-7803-4199-6 abstract page 171, right-hand column, line 30 -page 172, left-hand column, line 1 page 172, left-hand column, line 14-25 --- -/-</p>	1-3

Further documents are listed in the continuation of box C.

Patent family members are listed in annex.

* Special categories of cited documents

- *A* document defining the general state of the art which is not considered to be of particular relevance
- *E* earlier document but published on or after the international filing date
- *L* document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- *O* document referring to an oral disclosure, use, exhibition or other means
- *P* document published prior to the international filing date but later than the priority date claimed

- *T* later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- *X* document of particular relevance: the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- *Y* document of particular relevance, the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.
- *S* document member of the same patent family

Date of the actual completion of the international search

9 August 2001

Date of mailing of the international search report

17/08/2001

Name and mailing address of the ISA

European Patent Office, P B 5618 Patentlaan 2
 NL - 2280 HV Rijswijk
 Tel. (+31-70) 340-2040, Tx 31 651 epo nl.
 Fax (+31-70) 340-3016

Authorized officer

Dhondt, E

INTERNATIONAL SEARCH REPORT

International Application No
PCT/GB 01/00972

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	MCKEOWN N ET AL: "A quantitative comparison of iterative scheduling algorithms for input-queued switches" COMPUTER NETWORKS AND ISDN SYSTEMS, NL, NORTH HOLLAND PUBLISHING AMSTERDAM, vol. 30, no. 24, page 2309-2326 XP004150557 the whole document ---	1-3
A	PAO D C W ET AL: "LOOK-AHEAD RESERVATION-BASED SCHEDULING FOR INPUT-OUTPUT BUFFERED ATM SWITCH" IEE PROCEEDINGS: COMMUNICATIONS, GB, INSTITUTION OF ELECTRICAL ENGINEERS, vol. 145, no. 4, page 234-240 XP000784492 ISSN: 1350-2425 page 236, left-hand column, line 31-36 ----	1-3

THIS PAGE BLANK (USPTO)

